GDPS/PPRC
User Experiences

Ragnar Botnen

rabo01@handelsbanken.se

SHB overview

- Swedish bank
- About 10.000 employees
- About 600 branches in Sweden, Norway, Denmark, Finland, and UK
- Represented in Europe, Far East Asia, USA

Background

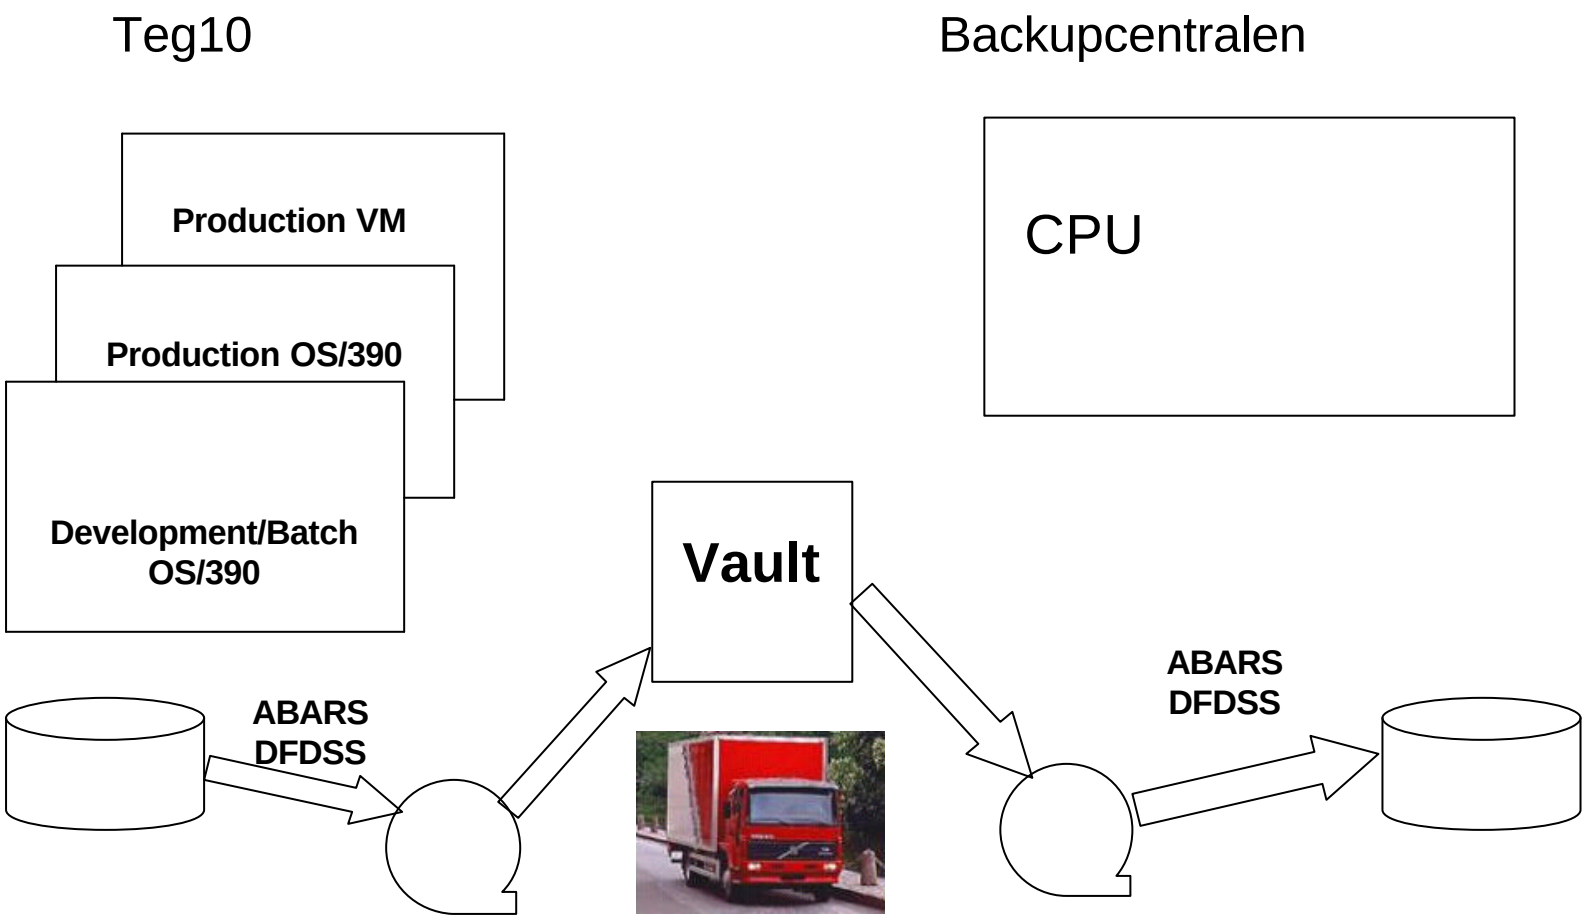
Teg10
Backupcentralen
Production VM
Production OS/390
Development/Batch
OS/390
CPU
Vault
ABARS
DFDSS
ABARS
DFDSS

Background

- **Disadvantages with old method**
  - RPO - Recovery Point Objective
    - Up to 24 hours of lost data
  - RTO – Recovery Time Objective
    - 48 – over 100 hours recovery
  - Difficult testing and maintenance
- **Goals**
  - RPO = 0
    - No data loss
  - RTO < 12 hours

Background

- Started 1995
  - Two datacenters
  - Mirrored data
    - Only production data mirrored
    - Separate production data from test data
  - Automation and Operation
    - Both sites operated from one place
  - Recovery procedures
    - Failure detection and alert
  - Maintenance procedures
    - Planned maintenance
  - GDPS

GDPS

- **Geographically Dispersed Parallel Sysplex**
  - **Set of disaster recovery and continuous availability solutions**
  - **Supports IBM's PPRC and XRC architectures**
  - **Failure recovery from a single point of control**
  - **Capability to manage your remote copy configuration and storage subsystem**
  - **Automation of Sysplex operational tasks**
  - **z/OS and Open System data**
  - **Open architecture**
  - **Application Independent**

GDPS

- **Prerequisites(GDPS/PPRC)**
  - **z/OS R4(e) or higher**
    - Not as guest under z/VM
  - **Netview**
    - The Enterprise or Procedural level of Tivoli NetView for OS/390 V1.4 (or higher)
    - The Enterprise or Procedural level of Tivoli NetView for z/OS V5.1 (or higher)
  - **REXX Runtime Library**
  - **System Automation**
    - System Automation for OS/390 V2.2 or higher
    - Tivoli System Automation for z/OS V2.3 or higher
    - Tivoli System Automation for z/OS V3.1 or higher
  - **The following hardware must exist in *both* sites**:
    - CPCs
    - Sysplex Timers
    - Disk subsystems
    - HMCs
  - **All disk subsystems that will be managed by GDPS must support the PPRC Freeze function. This function is also known as the CGroup Freeze/RUN architecture**.

GDPS-Terminology

- **PPRC**
  - **Metro Mirror (Synchronous PPRC)**
  - Global Mirror (Asynchronous PPRC)
  - Metro/Global Copy (Asynchronous Cascading PPRC)
  - Global Copy (PPRC Extended Distance)
- XRC
  - –z/OS Global Mirror (XRC)
  - –Metro/Global Mirror for zSeries (XRC and Synchronous PPRC)
- FLASHCOPY
- GDPS Controlling system
  - Recovery control residing in Site2

GDPS-Terminology

- **PPRC Freeze**
  - Secondary DASD consistency
  - Freeze Trigger (mirroring problem)
  - **Freeze and Stop**
    - Freeze secondary DASD and stop all systems after freeze trigger
  - **Freeze and Go**
    - Freeze secondary DASD and go all systems after freeze trigger
- **HYPERSWAP**
  - Transparent swap of primary DASD between sites
  - Swap trigger (primary DASD problem)
- **FAILOVER/FAILBACK**
  - Keep PPRC-Status after hyperswap

Implementation

- GDPS used in production since October 1997 for:
  - Planned shutdowns and IPL:S
    - Software maintenance
  - Planned site switches
    - Hardware maintenance
  - Managing PPRC-environment
    - New DASD
  - Sysplex Resource Management
    - Managing Couple Datasets
    - Managing Coupling Facilities
  - Disaster tests
    - Twice every year

Milestones

- Started on HDS7700E 4Q 1998
  - 500MB mirrored
- Mirroring of all applicable data 2Q 2000
  - Not system data(SYSRES,PAGE,SPOOL etc)
- HDS 9960 2Q 2002
- BCP Internal Interface replaced AOM 3Q 2002
- Hyperswap support 2Q 2004
  - **ALL** data mirrored
- CBU – Capacity Backup 4Q 2004

Milestones

- MFA – Message Flood Automation 2Q 2005
- HDS USP/Tagmastore 4Q 2005
  - Support for Failover/failback
- CF HINT GDPS 3.3 1Q 2006
  - Coupling Facility site awareness
- Full Parallel Sysplex with Data Sharing 2Q 2006
- xDR – GDPS Multiplatform Resiliency for zSeries 4Q 2006
  - Hyperswap support for z/VM and Linux for zSeries

Justification

- RTO < 2 hours
- RPO=0
- D/R tests much simpler and 100% successful
- Cost justified in 3 years
- No other solution known to us!

SHB IT Operations hardware

- Two centers
- 2 x 2094-705
  - 5 CP, 1 zAAP, 1 zIIP, 1 IFL, 2 ICF, 48GB real
  - 5 CP, 1 zAAP, 1 IFL, 2 ICF, 64GB real
- 16TB Hitachi USP DASD in each site
  - PPRC mirrored across 5km
  - 3 separated Single Mode Fiber Connections
  - FICON attached, fibre channel used for mirroring
- 2 STK silos in primary site, 1 in secondary
- About 12.000 tapes
  - 110TB in primary site

SHB IT Operations z/OS

- 2 Sysplexes, 5-way for production, 3-way for sysprog test
  - Sysprog test environment with dedicated SSID:s, CHPID:s and PPRC-links
- CF-duplexing used between sites for appropriate structures
- z/OS 1.8, IMS V9, DB2 V8, MQ 5.3.1, WAS 6
- GDPS/PPRC 3.3
- I/O 11.000/second at peak time. Approx 25-30% writes
- 200 IMS-trans/seconds at peak time
- 130 JAVA-servlets/seconds at peak time

SHB IT Operations z/OS

- Internet bank runs on WAS 6 in z/OS
  - One ND Cell with 7 App servers
  - Presentation layer only – no business logic (yet)
  - Sysplex load sharing across two images on separate sites
- Swedish branch office presentation layer also on WAS 6
  - One ND Cell with 4 App servers per z/OS image
  - Sysplex load sharing across two images on separate sites
- Business logic runs in IMS + DB2
  - IMS/DB2 shared across two z/OS images on separate sites
- Batch load shared across two z/OS images on separate sites
- TCP/IP shared across two+ z/OS images on separate sites

SHB IT Operations z/VM

- z/VM 5.1 One LPAR for production and one for sysprog test(xDR)
  - Production environment with dedicated SSID:s and PPRC-links
  - xDR test environment with dedicated SSID:s, CHPID:s and PPRC-links
- z/VM
  - Financial applications
  - Moved to secondary site if primary CPU goes down
- Approx 20 Linux images runs as guests under z/VM
  - File and print serving
  - New workload
    - Tivoli Workload Scheduler
    - Tivoli Framework
    - Tivoli Gateway
    - …and more
  - Dedicated Linux image for xDR proxy SuSE sles 9

GDPS Test environment

- Simulated DASD failure at primary site(Unplanned Hyperswap)
  - V OFFLINE,xxxx,FORCE of mirrored DASD or
  - V PATH(xxxx,yy),OFF and then pulling last channel for mirrored DASD
- Simulated DASD and system failure at primary site(Unplanned Hyperswap)
  - V OFFLINE,xxxx,FORCE of mirrored DASD or
  - V PATH(xxxx,yy),OFF and then pulling last channel for mirrored DASD
  - System reset from outside GDPS(HMC)
- Simulated link failure(Freeze and Stop)
  - CSUSPEND outside of GDPS
  - Policy used at SHB - ”SWAP,STOP”
- Testing after all major software and hardware upgrades

GDPS Disaster Recovery Verification

- Same failure simulation methods as in testplex
  - Verification done twice every year
  - All verification done in real production environment
  - Simulated DASD failure without impact for end-users
  - Simulated link failure stops all applications – big impact!
  - Application owners participate
  - Verification coordinated with other platforms

Site1 Disaster

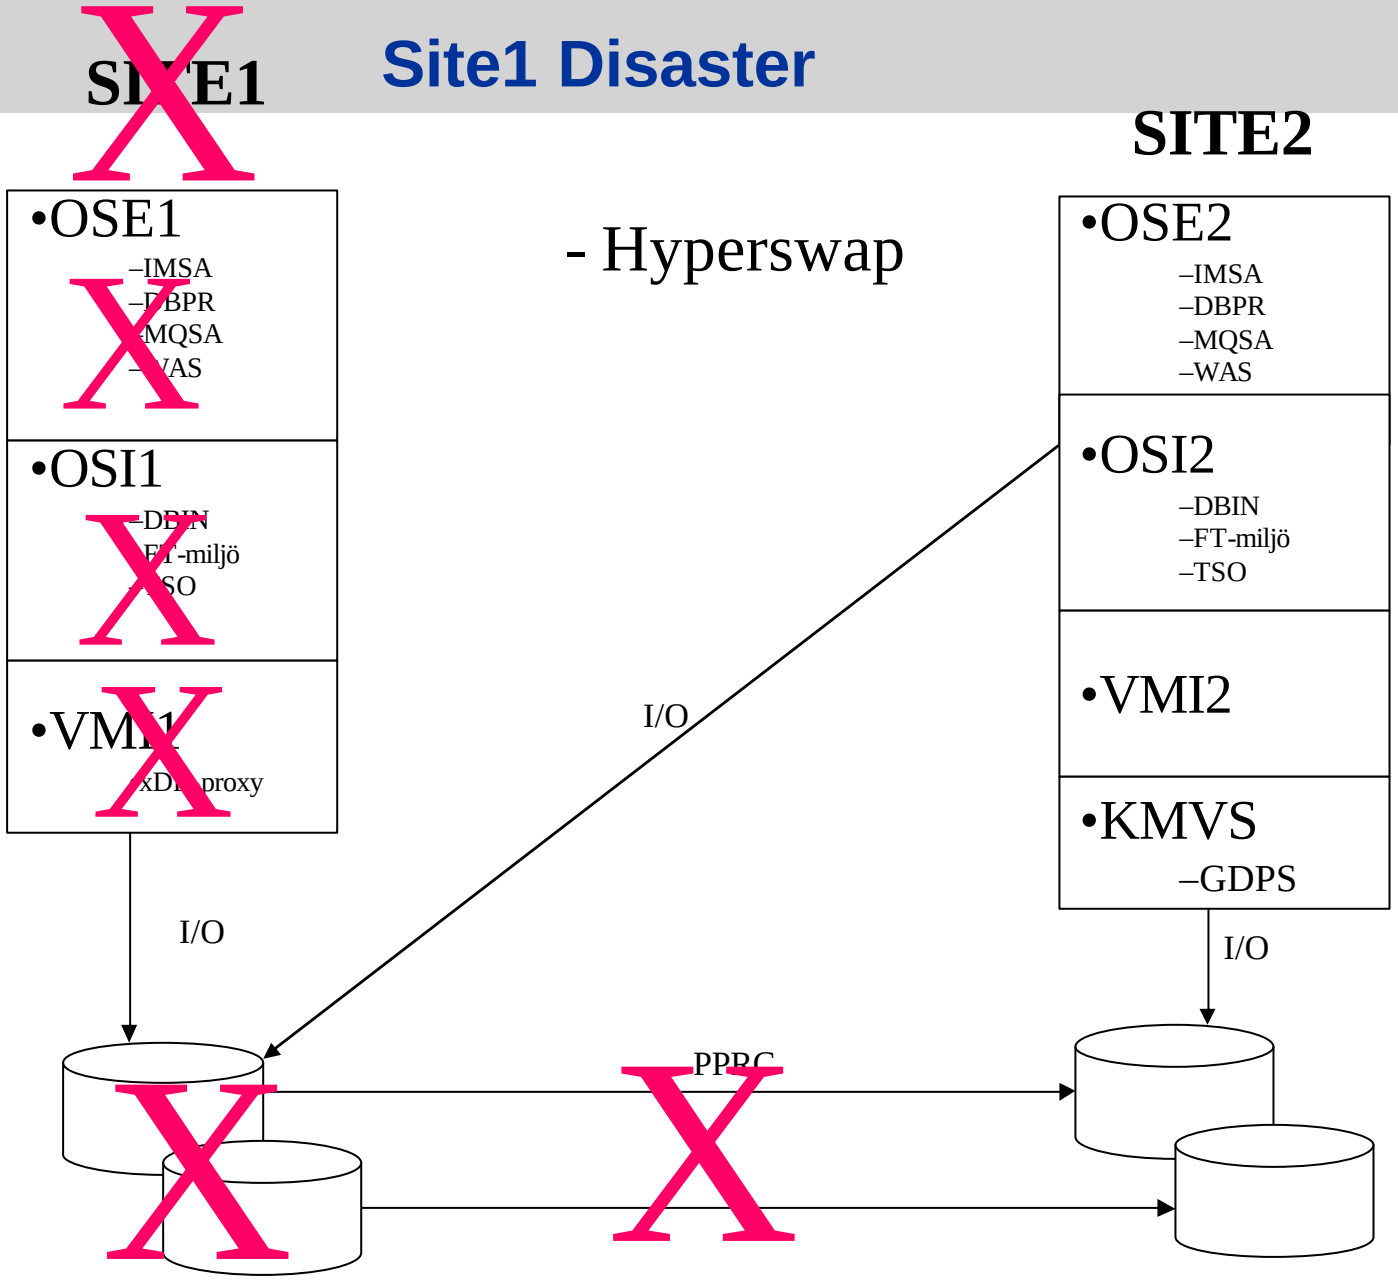

Site1 Disaster

- Hyperswap

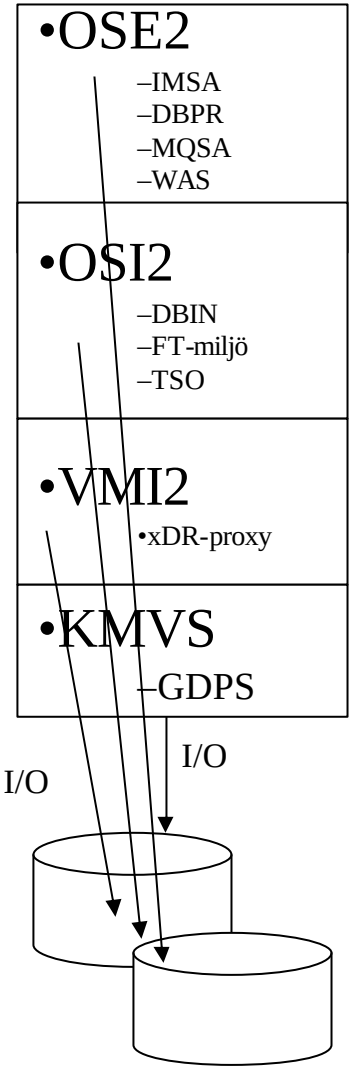

Future

- OPEN LUN for iSeries?
- Metro/Global Copy (Asynchronous Cascading PPRC)?
  - 3 Site solution
- 2 GDPS controlling systems?
- OPEN LUN for other platforms?
- I/O Timing triggers?

Questions?